Research Article

DRIVEN BY doi®

High speed fir filters using add and shift method based on FPGA

Ch.Venkateswarlu [1] and T.Anilkumar [2]

Corresponding Author:
tvakumar2000@yahoo.co.in

ABSTRACT

In this paper, a method for implementing high speed Finite Impulse Response (FIR) filters using just registered adders and hardwired shifts is presented. A modified common subexpression elimination algorithm is extensively used to reduce the number of adders. The target is on optimizations to Xilinx Virtex II devices where the implementations are compared with those produced by Xilinx CoregenTM using Distributed Arithmetic. It is observed that up to 50% reduction in the number of slices and up to 75% reduction in the number of LUTs for fully parallel implementations and also observed up to 50% reduction in the total dynamic power consumption of the filters. The designs implemented in this method perform significantly faster than the MAC filters, which use embedded multipliers.

Keywords: Combinational Logic Blocks, Finite Impulse Response Filter, Multiply and Accumulate, Look up Table.

[1] Department of ECE, Vaageswari Engineering College, Karimnagar, India 505 481.
[2] Department of ECE, SR Engineering College, Warangal, India 506 371.
.

High speed fir filters using add and shift method based on FPGA

Ch.Venkateswarlu [1] and T.Anilkumar [2]

[1]Department of ECE, Vaageswari Engineering College, Karimnagar, India 505 481.
[2]Department of ECE, SR Engineering College, Warangal, India 506 371.
[1]ecehodum@gmail.com and [2]tvakumar2000@yahoo.co.in

ABSTRACT

In this paper, a method for implementing high speed Finite Impulse Response (FIR) filters using just registered adders and hardwired shifts is presented. A modified common subexpression elimination algorithm is extensively used to reduce the number of adders. The target is on optimizations to Xilinx Virtex II devices where the implementations are compared with those produced by Xilinx CoregenTM using Distributed Arithmetic. It is observed that up to 50% reduction in the number of slices and up to 75% reduction in the number of LUTs for fully parallel implementations and also observed up to 50% reduction in the total dynamic power consumption of the filters. The designs implemented in this method perform significantly faster than the MAC filters, which use embedded multipliers.

Keywords: Combinational Logic Blocks, Finite Impulse Response Filter, Multiply And Accumulate, Look Up Table.

I. INTRODUCTION

FPGAs are being increasingly used for a variety of computationally intensive applications, mainly in the realm of Digital Signal Processing (DSP) and communications. Due to rapid increases in the technology, current generation of FPGAs contain a very high number of Configurable Logic Blocks (CLBs), and are becoming more feasible for implementing a wide range of applications. The high non-recurring engineering (NRE) costs and long development time for ASICs are making FPGAs more attractive for application specific DSP solutions. DSP functions such as FIR filters and transforms are used in a number of applications such as communication and multimedia. These functions are major determinants of the performance and power consumption of the whole system. Therefore it is important to have good tools for optimizing these functions.

Eq (1) represents the output of an L tap FIR filter, which is the convolution of the latest L input samples. L is the number of coefficients h (k) of the filter, and x (n) represents the input time series.

$$y[n] = \sum h[k] \, x \, [n-k] \qquad k = 0, \, 1 \ldots L-1 \qquad (1)$$

The conventional tapped delay line realization of this inner product is shown in Figure 1. This implementation translates to L multiplications and L-1 additions per sample to compute the result. This can be implemented using a single multiply and Accumulate (MAC) engine, but it would require L MAC cycles, before the next input sample can be processed. Using a parallel implementation with L MACs can speed up the performance L times. A general purpose multiplier occupies a large area on FPGAs. Since all the multiplications are with constants, the full flexibility of a general purpose multiplier is not required, and the area can be vastly reduced using techniques developed for constant multiplication. Though most of the current generation FPGAs such as Virtex II™ have embedded multipliers to handle these multiplications, the number of these multipliers is typically limited. Furthermore, the size of these multipliers is limited to only 18 bits, which limits the precision of the computations for high speed requirements. The ideal implementation would involve a sharing of the Combinational Logic Blocks (CLBs) and these multipliers. In this paper, a technique that is better than conventional techniques for implementation on the CLBs is presented.

Fig. 1. A MAC FIR filter block diagram.

An alternative to the above approach is Distributed Arithmetic (DA) which is a well known method to save resources. Using DA method, the filter can be implemented either in bit serial or fully parallel mode to trade bandwidth for area utilization. Assuming coefficients c[n] are known constants, equation (1) can be rewritten as follows:

$$y[n] = \sum c[n] \cdot x[n] \qquad n = 0, 1, \ldots, N-1 \qquad (2)$$

x[n] can be represented by:

$$x[n] = \sum x_b[n] \cdot 2^b \qquad b = 0, 1, \ldots, B-1 \qquad (3)$$
$$x_b[n] \in [0, 1]$$

Where $x_b[n]$ is the b^{th} bit of x[n] and B is the input width. Finally, the inner product can be rewritten as follows:

$$y = \sum c[n] \sum x_b[k] \cdot 2^b$$
$$= c[0] (x_{B-1}[0]2^{B-1} + x_{B-2}[0]2^{B-2} + \ldots + x_0[0]2^0) + c[1] (x_{B-1}[1]2^{B-1} + x_{B-2}[1]2^{B-2} + \ldots + x_0[1]2^0) + \ldots + c[N-1] (x_{B-1}[N-1]2^{B-1} + x_{B-2}[0]2^{B-2} + \ldots + x_0[N-1]2^0)$$
$$= (c[0] x_{B-1}[0] + c[1] x_{B-1}[1] + \ldots + c[N-1] x_{B-1}[N-1])2^{B-1} + (c[0] x_{B-2}[0] + c[1] x_{B-2}[1] + \ldots + c[N-1] x_{B-2}[N-1])2^{B-2}$$
$$+ \ldots$$
$$+ (c[0] x_0[0] + c[1] x_0[1] + \ldots + c[N-1] x_0[N-1])2^0$$
$$= \sum 2^b \sum c[n] \cdot x_b[k] \qquad (4)$$

Where n=0, 1… N-1 and b=0, 1… B-1
The coefficients in most of DSP applications for the multiply accumulate operation are constants. The partial products are obtained by multiplying the coefficients c_i by multiplying one bit of data x_i at a time in AND operation. These partial products should be added and the result depends only on the outputs of the input shift registers. The AND functions and adders can be replaced by Look up Tables (LUTs) that gives the partial product. This is shown in Figure 2. Input sequence is fed into the

shift register at the input sample rate. The serial output is presented to the RAM based register stores the data in a particular address. The outputs of registered LUTs are added and loaded to the scaling accumulator from LSB to MSB and the result which is the filter output will be accumulated over the time. For an n bit input, n+1 clock cycles are needed for a symmetrical filter to generate the output. Shift registers (registers are not shown in Figure for simplicity) at the bit clock rate which is n+1 times (n is number of bits in a data input sample) the sample rate. The RAM based shift

In conventional MAC method with a limited number of MAC engines, as the filter length is increased, the system sample rate is decreased. This is not the case with serial DA architectures since the filter sample rate is decoupled from the filter length. As the filter length is increased, the throughput is maintained but more logic resources are consumed.

Though the serial DA architecture is efficient by construction, its performance is limited by the fact that the next input sample can be processed only after every bit of the current input samples are processed. Each bit of the current input samples takes one clock cycle to process.

A popular technique for implementing the transposed form of FIR filters is the use of a multiplier block, instead of using multipliers for each constant as shown in Figure 4. The multiplications with the set of constants {h_k} are replaced by an optimized set of additions and shift operations, involving computation sharing. Further optimization can be done by factorizing the expression and finding common sub expressions. The performance of this filter architecture is limited by the latency of the biggest adder and is the same as that of the PDA.

The main contribution in this paper is the development of a novel algorithm for optimizing the multiplier block for FIR filters, using a modified algorithm for common sub expression elimination. The goal of the algorithm is to produce a filter that can provide the maximum sample rate with the least amount of hardware. Our algorithm takes into account the specific features of FPGA slices to reduce the total number of occupied slices. The reduced number of slices also leads to a reduction in the total power on the FPGA.

We compare our results with the industry standard Xilinx Coregen™, where we compare the total area and power consumption.

The rest of the paper is organized as follows: Section 2 presents some related work. In Section3, we describe our filter architecture. In Section 4, we present our optimization algorithm for reducing the total area of the design. In Section 5, we describe our experimental setup and present our results. Finally we conclude the paper in Section 6.

II. RELATED WORK

Multiplications with constants have to be performed in many signal processing and communication applications such as FIR filters, audio, video and image processing. Since implementing a general purpose multiplier is expensive on an FPGA and since we do not really need such a multiplier, when one of the operands is a constant, there has been a lot of work on deriving efficient structures for constant multiplications. All these techniques are based on computing constant multiplications using table lookups and additions. The method of Distributed Arithmetic which is the most popular method for implementing Multiplier less FIR filters, is also based on table lookup. The Xilinx™ CORE Generator has a highly parameterizable, optimized filter core for implementing digital FIR filters. Based on both Distributed Arithmetic as well as MAC (Multiply Accumulate) based architectures. It generates synthesized core that targeting a wide range of Xilinx devices. The MAC based implementations make use of the embedded DSP slices on the FPGA devices. In this work, we primarily compare our technique with the Coregen implementation of the Distributed Arithmetic, since that also is a multiplier less technique. We show that our designs are much more area efficient than the DA based approach for fully parallel filters. We also compare our method with MAC based implementations, where we achieve significantly higher performance

Though there has been a lot of work on optimizing constant multiplications using adders and employing redundancy elimination, they have not been effectively used for FIR filter design. The closest work to implementing filters with adders is in, FIR filters are implemented using the Add and Shift method. Canonical Signed Digit (CSD) encoding is used for the coefficients to minimize the number of additions. The paper discusses how high speed implementations can be achieved by registering each

adder, due to which the critical path becomes equal to the delay of the adder. Registering an adder output comes at no extra cost on an FPGA because of the presence of a D flip flop at the output of each LUT. In comparison with, we extensively use common subexpression elimination for reducing the number of adders and therefore area. Furthermore, our designs can run with sample rates as high as 252 Msps (Million samples per second), whereas the designs in can run only at 78.6 Msps.

In comparison with the other algorithms for common subexpression elimination, our method takes into account the structure of the FPGA slices and takes into account both the cost of adders and registers when performing the optimization. Furthermore, we provide comprehensive evidence of the benefits of our technique through experimental results, where we compare our results with those produced by industry standard tools

.

III. FILTER ARCHITECTURE

We base our filter architecture on the transposed form of the FIR filter as shown in Figure 1. The filter can be divided into two main parts, the multiplier block and the delay block, and is illustrated in Figure 4. In the multiplier block, the current input variable x[n] is multiplied by all the coefficients of the filter to produce the y_i outputs. These y_i outputs are then ^{trees} delayed and added in the delay block to produce the filter output y[n].

We perform all our optimizations in the multiplier block. The constant multiplications are decomposed into registered additions and hardwire shifts. The additions are performed using two input adders, which are arranged in the fastest tree structure. We use registered adders, so that the performance of the filter is only limited by the slowest adder. We use common subexpression elimination extensively, to reduce the number of adders, which leads to a reduction in the area. To synchronize all the intermediate values in the computation, we insert registers in the dataflow, where


Fig. 3. Registered Adder at no additional cost.

Performing subexpression elimination can sometimes increase the number of registers substantially, and the overall area could possibly increase. Consider the two expressions F_1 and F_2 which could be part of the multiplier block.

$F_1 = A + B + C + D$

$F_2 = A + B + C + E$

Both the expressions have a minimum critical path of two addition cycles. These expressions require a total of six registered adders for the fastest implementation, and no extra registers are required. From the expressions we can see that the computation $A + B + C$ is common to both the expressions. Since both D and E need to wait for two addition cycles to be added to $(A + B + C)$, we need to use two registers each for D and E, such that new values for A,B,C,D and E can be read in at each clock cycle. A more careful sub expression elimination algorithm would only extract the common sub expression $A + B$ (or $A+C$ or $B + C$). The number of adders is decreased by one from the original, and no additional registers are added. This is illustrated in Figure 8. The algorithm for performing this kind of optimization is described in the next section.


Fig. 4. Unoptimized expression.

Fig. 5. Extracting common subexpression (A+B).

IV. OPTIMIZATION ALGORITHM

The goal of our optimization is to reduce the area of the multiplier block by reducing the number of adders and any additional registers required for the fastest implementation of the FIR filter. We first give a brief overview of the common sub expression elimination methods. A detailed description can be found in [22]. We then present the modified optimization algorithm to be used for our work.

A) Overview of common subexpression elimination

We use a polynomial transformation of constant multiplications. Given a representation for the constant C, and the variable X, the multiplication C*X can be represented as a summation of terms denoting the decomposition of the multiplication into shifts and additions as

$$C * X = \sum_i \pm XL^i \qquad (5)$$

The terms can be either positive or negative when the constants are represented using signed digit representations such as the Canonical Signed Digit (CSD) representation. The exponent of L represents the magnitude of the left shift and the i's represent the digit positions of the non-zero digits of the constants. For example the multiplication $7*X = (100-1)$ CSD$*X = X{<}{<}3 - X = XL^3 - X$, using the polynomial transformation.

We use the divisors to represent all possible common sub expressions. Divisors are obtained from an expression by looking at every pair of terms in the expression and dividing the terms by the minimum exponent of L. For example in the expression $F = XL^2 + XL^3 + XL^5$, consider the pair of terms $(+XL^2 + XL^3)$. The minimum exponent of L in the two terms is L^2. Dividing by L^2, we get the divisor $(X + XL)$. From the other two pairs of terms $(XL^2 + XL^5)$ and $(XL^3 + XL^5)$, we get the divisors $(X + XL^3)$ and $(X + XL^2)$ respectively. These divisors are significant, because every common subexpression in the set of expressions can be detected by performing intersections among the set of divisors.

B) Optimization algorithm

We first calculate the minimum number of registers required for our design. We calculate this by arranging the original expressions in the fastest possible tree structure, and then inserting registers. For example, for the six term expression $F = A + B + C + D + E + F$, we have the fastest tree structure with three addition steps, and we require one register to synchronize the intermediate values, such that new values for A, B, C, D, E, F can be read in every clock cycle. This is illustrated in Fig. 9.


Fig. 6. Calculating registers required for fastest evaluation.

Consider the expressions shown in Fig. 6. We need six registered adders and no additional registers for the fastest evaluation of F_1 and F_2. Now consider the selection of the divisor $d_1 = (A+B)$. This divisor saves one addition and does not increase the number of registers. Divisors $(A + C)$ and $(B + C)$ also have the same value, but $(A+B)$ is selected randomly. The expressions are now rewritten as:

$d_1 = (A + B)$, $F_1=d_1+C+D$ & $F_1=d_1+C+E$

Optimization algorithm to reduce area:

Reduce Area ({Pi})

{

{Pi} = Set of expressions in polynomial form; {D} = Set of divisors = ϕ;

//Step 1: Creating divisors and calculating minimum number of registers required

For each expression P_i in {Pi}

{

{D_{new}} = Find Divisors (P_i);

Update frequency statistics of divisors in {D};

{D} = {D} ∪ {Dnew};

P_i->MinRegisters = Calculate Minimum registers required for fastest evaluation of P_i;

}

//Step 2: Iterative selection and elimination of best divisor while (1)

{

Find d = Divisor in {D} with greatest Value;

// Value = Num Additions reduced – Num Registers Added;

If (d == NULL) break;

Rewrite affected expressions in {Pi} using d;

Remove divisors in {D} that have become invalid;

Update frequency statistics of affected divisors;

{D_{new}} = Set of new divisors from new terms added by division;

{D} = {D} ∪ {D_{new}};

}

}

After rewriting the expressions and forming new divisors, the divisor $d_2 = (d_1 + C)$ is considered. This divisor saves one adder, but introduces five additional registers, as can be seen in Figure 7. Therefore this divisor has a value of - 4. No other valuable divisors can be found and the iteration stops. We end up with the expressions shown in the algorithm.

Table 1. Filter Synthesis using Add Shift method.

Filter (# taps)	Slices	LUTs	FFs	Performance (Msps)
6	524	774	1012	245
10	781	1103	1480	222
13	929	1311	1775	199
20	1191	1631	2288	199
28	1774	2544	3381	199
41	2475	3642	4748	222
61	3528	5335	6812	199
119	6484	9754	12539	205
151	8274	12525	15988	199

V. EXPERIMENTS

The goal of our experiments was to compare the number of resources consumed by add and shift method with that produced by the cores generated by the commercial Coregen™ tool, based on Distributed Arithmetic. Besides the resources, we also compared the power consumption of the two implementations, and also measured the performance. For our experiments, we considered 9 FIR filters of various sizes. We targeted the Xilinx Virtex II device for our experiments. The constants were normalized to 17 digit of precision and the input samples were assumed to be 12 bits wide.

Table 2. Filter Synthesis using Coregen (PDA method).

Filter (# taps)	Slices	LUTs	FFs	Performance (Msps)
6	264	213	509	251
10	474	406	916	222
13	386	334	749	252
20	856	705	1650	250
28	1294	1145	2508	227
41	2154	1719	4161	223
61	3264	2591	6303	192
119	6009	4821	11551	203
151	7579	6098	14611	180

From the results we can observe up to 50% reduction in dynamic power consumption. We did not include the quiescent power into our calculation since that value is the same for both methods. The power consumption is the result of applying the same test stimulus to both designs and measuring the power using XPower tools provided by Xilinx ISE software.


Fig. 7. Comparison of Reduction in Resourses for considered algorithms.

Fig. 8. Comparison of dynamic power consumption for considered algorithms.

Comparison with MAC filters using embedded multipliers Coregen ™ can produce FIR filters based on the Multiply Accumulate (MAC) method, which makes use of the embedded multipliers and DSP blocks. We implemented the FIR filters using the MAC method to compare the resource usage and performance with add and shift method. Due to tool limitations we had to do the experiments for Virtex IV device. We present the synthesis results in terms of number of slices on the Virtex IV device and the performance in Msps in Table 3.

Table 3. Comparing with MAC filter on VirtexIV.

Filter (# taps)	Add Shift Method		MAC Filter	
	Slices	Msps	Slices	Msps
6	264	296	219	262
10	475	296	418	253
13	387	296	462	253
20	851	271	790	251
28	1303	305	886	251
41	2178	296	1660	243
61	3284	247	1947	242
119	6025	294	3581	241
151	7623	294	7631	215

VII. CONCLUDING REMARKS

In this paper we presented a multiplier less technique, based on add and shift method and common subexpression elimination for low area, low power and high speed implementations of FIR filters. We validated our techniques on Virtex II™ devices where we observed significant area and power reductions over traditional Distributed Arithmetic based techniques. In future, we would like to modify our algorithm to make use of the limited number of embedded multipliers available on the FPGA devices.

REFERENCES

[1]   K.D.Underwood and K.S.Hemmert, "Closing the Gap: CPU and FPGA Trends in Sustainable Floating-Point BLAS Performance," presented at Intl. Symp. On Field-Programmable Custom Computing Machines, California, USA, 2004.

[2]   L.Zhuo and V.K.Prasanna, "Sparse Matrix-Vector Multiplication on FPGAs," presented at Intl. Symp. On Field Programmable Gate Arrays (FPGA), Monterey, CA, 2005.

[3]   Y.Meng, A.P.Brown, R.A.Iltis, T.Sherwood, H.Lee, and R.Kastner, "MP Core: Algorithm and Design Techniques for Efficient Channel Estimation in Wireless Applications," presented at Design Automation Conference (DAC), Anaheim, CA, 2005.

[4]   B. L. Hutchings and B. E. Nelson, "Gigaop DSP on FPGA," presented at Acoustics, Speech, and Signal Processing, 2001. Proceedings. (ICASSP '01). 2001 IEEE Intl. Conf. on, 2001.

[5]   A.Alsolaim, J.Becker, M.Glesner, and J.Starzyk, "Architecture and Application of a Dynamically Reconfigurable Hardware Array for Future Mobile Communication Systems," presented at Intl. Symp. On Field Programmable Custom Computing Machines (FCCM), 2000.

[6]   FPGA," presented at Intl. Symp. On Field-Programmable Custom Computing Machines (FCCM), 2002.